arXiv:1412.1253v1 [math.NA] 3 Dec 2014

PRECONDITIONERS FOR HIERARCHICAL MATRICES BASED ON THEIR EXTENDED SPARSE FORM

DARIA SUSHNIKOVA[¶] AND IVAN V. OSELEDETS[‡¶]

Abstract. In this paper we consider linear systems with dense-matrices which arise from numerical solution of boundary integral equations. Such matrices can be well-approximated with $\mathcal{H}^2$-matrices. We propose several new preconditioners for such matrices that are based on the equivalent *sparse extended form* of $\mathcal{H}^2$-matrices. In the numerical experiments we show that the most efficient approach is based on the so-called reverse-Schur preconditioning technique.

Key words.
$\mathcal{H}^2$-matrix, integral equations, preconditioning

AMS subject classifications.

1. Introduction. Many physical problems in acoustics, electrostatics [21] and some other [2] areas lead to boundary and volume integral equations with non-local kernels. Numerical solution of integral equations is challenging, since even the computation of all the matrix elements is offen not possible for the problems of interest. Fortunately, the matrices, arising from the discretization of integral equations can be approximated well with *data-sparse matrices*. Among the most important approaches are hierarchical matrices ($\mathcal{H}$-matrices) [5], mosaic-skeleton method [23] , and hierarchical semiseparable matrices (HSS-matrices) [22, 16, 7, 8]. All of those classes of matrices correspond to the idea of block low-rank approximation and have their roots in the classical fast multipole method (FMM) [11, 12, 19]. In this paper we consider linear systems with $\mathcal{H}^2$-matrices [13, 4]. An $\mathcal{H}^2$-matrix can be multiplied by a vector fast, thus iterative methods can be used to solve linear systems. But for efficient solution of a linear system it is not enough. If matrix is ill-conditioned we have to use preconditioners. A very efficient approach is based on the approximate factorization of hiearchical matrices [3]. Algorithms with almost linear complexity have been proposed and successfully applied to many problems. The disadvantage of these methods is that the prefactor can be quite large, and the memory required to store the L and U factors can also be large. In the recent years several approaches have been proposed for fast direct methods with HSS matrices. HSS matrix is equivalent to the $\mathcal{H}^2$-matrix corresponding to one-dimensional geometry, thus this structure is not fully suited for solving surface and volume integral equations, i.e. it can not give optimal complexity. Nevertheless, the actual computing times can be amazing [15, 10, 24].

In this paper we use the observation that the classical three-step matrix-by-vector product procedure can be rewritten as a big linear system (which we call sparse extended form or simply SE-form of the $\mathcal{H}^2$-matrix). Very similar ideas were presented in [6, 15, 2, 9]. In paper [1] SE form is used for building direct solver for system with $\mathcal{H}^2$-matrix. We propose a number of new methods for preconditioning systems with H2 Matrix, based on idea of SE-form. For small problem sizes (say, $N \sim 20\,000$) this gives an easy-to-implement approach for the solution of a given linear system with an $\mathcal{H}^2$-matrix, we have found that the memory requirements and the computational cost grow very fast. Therefore we propose several alternatives, which use SE-form as an

[‡]Skolkovo Institute of Science and Technology, Novaya St. 100, Skolkovo, Odintsovsky district, 143025 Moscow Region, Russia (i.oseledets@skolkovotech.ru)

[¶]Institute of Numerical Mathematics Russian Academy of Sciences, Gubkina St. 8, 119333 Moscow, Russia

auxiliary step for the creation of efficient preconditioners for the initial matrix. We numerically confirm the effectiveness of the proposed preconditioners on two surface integral equations. The code is available online as a part of the open-source Python package h2tools [17].

2. Notations and basic facts. In this section we recall basic definitions and notations for working with $\mathcal{H}^2$-matrices. This material is not new and can be found, for example in [4]. Let us consider a matrix $A \in \mathbb{R}^{m\times n}$. First, we will need several definitions:

DEFINITION 2.1. *Block cluster tree $\mathcal{T}_r$ ($\mathcal{T}_c$) of rows (columns) of matrix A is a tree where:*

1. *Each node $\mathcal{T}_r^t$ ($\mathcal{T}_c^s$) is associated with some group of rows (columns)*
2. *Root node $\mathcal{T}_r^0$ ($\mathcal{T}_c^0$) contains all rows (columns)*
3. *If some group of rows (columns) is divided into subgroups, then the corresponding node has child nodes associated with those subgroups.*

DEFINITION 2.2. *Let $\mathcal{T}_c$ and $\mathcal{T}_r$ be block cluster trees of columns and rows of the matrix A. Each pair $p = (v, w)$ where $v \in \mathcal{T}_r, w \in \mathcal{T}_c$ represents some block in matrix A. We assume that there is some rule that divides the set of blocks into two classes, namely* close blocks *and* far blocks*. The nodes of trees $i \in \mathcal{T}_c$ and $j \in \mathcal{T}_r$ are close (far) if the pair $p = (i, j)$ is close (far). If some node j of one tree is close to all children of some node i of another tree then the node j is close to the node i. Denote by $\mathcal{P}_{close}$ ($\mathcal{P}_{far}$) the set of all close (far) pairs $p = (i, j), i \in \mathcal{T}_r, j \in \mathcal{T}_c$*

DEFINITION 2.3 (Cluster basis.). *[4, p. 54] Let $K = (K_i)_{i\in\mathcal{T}_I}$ be a family of finite index sets. Let $R = (R_i)_{i\in\mathcal{T}_r}$ be a family of matrices satisfying $R_i \in \mathbb{R}^{I\times K_t}$ for all $i \in \mathcal{T}_r$. Then R is called a cluster basis with rank distribution K and the matrices R_i are called cluster basis matrices.* The $\mathcal{H}^2$-matrix structure is related to the *nestedness property.*

DEFINITION 2.4 (Nested cluster basis.). *[4, p. 54] Let R be a cluster basis with rank distribution K. R is called nested if there is a family $(E_i)_{i\in\mathcal{T}_r}$ of matrices satisfying the following conditions:*

1. *For all $i \in \mathcal{T}_r$ and all $i' \in \operatorname{sons}(i)$, we have $E_{i'} \in \mathbb{R}^{K_{i'}\times K_i}$.*
2. *For all $i \in \mathcal{T}_r$ with $\operatorname{sons}(i) \neq \varnothing$, the following equation holds:*

$$R_i = \sum_{i'\in sons(i)} R_{i'} E_{i'}, \tag{2.1}$$

The matrices E_t are called transfer matrices or expansion matrices. We can consider matrix A as a sum of two matrices: $A = C + F$, where C is constructed from close blocks $p \in \mathcal{P}_{close}$ and F is constructed from far blocks $p \in \mathcal{P}_{far}$. Now we can define an $\mathcal{H}^2$-matrix:

DEFINITION 2.5 ($\mathcal{H}^2 - representation$). *[4, p. 56] Let $A \in \mathbb{R}^{I\times J}$. Let $\mathcal{T}_c$ and $\mathcal{T}_r$ be block cluster trees of columns and rows of matrix A. R is a nested cluster basis of rows and L is a nested cluster basis of columns. If there is a family $S = (S_p)_{p\in\mathcal{P}_{far}}$ of matrices , for all $p = (i, j) \in \mathcal{P}_{far}$ then A is an $\mathcal{H}^2$-matrix and*

$$A = F + C = \sum_{(i,j)\in\mathcal{P}_{far}} R_i S_p L_j^* + \sum_{(i,j)\in\mathcal{P}_{close}} C_p \tag{2.2}$$

Figure 2.1 shows close blocks (black), far blocks (white) and an example of block row (shaded). The block row consists of columns of the blocks that are separated from the node i.

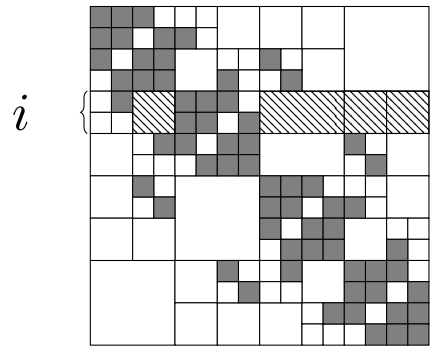

Fig. 2.1: Illustration of a matrix with $\mathcal{H}^2$-structure.

The construction of cluster trees, block cluster nested trees, could be done using the standard procedure, see, for example, [20]. The crucial task here is to compute the cluster basis. In our numerical experiments we have used the method proposed in [18], however other techniques maybe used as well. Summarizing, an $\mathcal{H}^2$-matrix A is defined by cluster trees $\mathcal{T}_c$ and $\mathcal{T}_r$, and the following sets of matrices $R = (R_i),\ i \in \mathcal{T}_r$, $C = (C_p),\ p = (i,j) \in \mathbb{P}_{close},\ i \in \mathcal{T}_r,\ j \in \mathcal{T}_c$, $S = (S_p),\ p = (i,j) \in \mathbb{P}_{far}, i \in \mathcal{T}_r,\ j \in \mathcal{T}_c$, $L = (L_j),\ j \in \mathcal{T}_c$,

2.1. Matrix-vector multiplication. Matrix-by-vector multiplication algorithm for the $\mathcal{H}^2$-matrix [4, p. 59-63] plays the key role in this paper. Its formal description is given in Algorithm 3. It is convenient to distribute the vector x over the nodes of the row cluster tree: $x = (x_i)\ i \in \text{leaves}(\mathcal{T}_r)$. The resulting vector $y = Ax$ is also computed in the form $y = (y_j)\ j \in \text{leaves}(\mathcal{T}_c)$. Note that in the algorithm we use there is an additional operation that transfers x to the leaves of the tree $\mathcal{T}_r$ with the help of the matrices $D = (D_i),\ i \in leaves(\mathcal{T}_r)$, and $E = (E_j),\ j \in leaves(\mathcal{T}_c)$.

Algorithm 1: Forward transformation

```
Procedure ForwardTransformation(i, R, x̂)
    if sons(i) ≠ 0 then
        x̂_{father(i)} := R_i x̂_i
    else
        x̂_i := 0
        for j ∈ sons(i) do
            x̂_j := ForwardTransformation(j, R, x̂)
            x̂_i := x̂_i + R_j x_j
    return: x̂
```

Algorithm 2: Backward transformation

Procedure `BackwardTransformation`$(i, L, \hat{y})$
  if sons$(i) \neq 0$ **then**
    $\hat{y}_{father(i)} := L_i \hat{y}_i$
  else
    for $j \in$ sons(i) **do**
      $\hat{y}_j := \hat{y}_j + L y_i$
      $\hat{y} :=$ `BackwardTransformation`$(j, L, \hat{y})$
  `return:`$(\hat{y})$

Algorithm 3: Matrix-vector multiplication.

input : $\mathcal{H}^2$-matrix A = $\{\mathcal{T}_r, \mathcal{T}_c$, D, R, C, S, L, E$\}$, vector x
output: vector y
for $i \in \mathcal{T}_r$ **do**
  $\hat{x}_i := 0$
for $j \in \mathcal{T}_c$ **do**
  $\hat{y}_j := 0$
Step 1:
for $i \in \mathcal{T}_r$, sons$(i) = 0$ **do**
  `/* For all leaf nodes */`
  $\hat{x}_i := D_i x_i$
Step 2:
$\hat{x} :=$ `ForwardTransformation`$(root(\mathcal{T}_r), R, \hat{x})$
Step 3:
for $i \in \mathcal{T}_r$ **do**
  for $j \in \mathcal{T}_c$ **do**
    if $(i, j) \in \mathcal{P}_{far}$ **then**
      $\hat{y}_j := S_{(ij)} \hat{x}_i$
Step 4:
$\hat{y} :=$ `BackwardTransformation`$(root(\mathcal{T}_c), L, \hat{y})$
Step 5:
for $i \in \mathcal{T}_c$, sons$(i) = 0$ **do**
  $y_i := E_i \hat{y}_i$

On Figure 2.2 we give a graphical illustration of the matrix-vector product procedure. The complexity is $\mathcal{O}(N)$ and storage complexity is also $\mathcal{O}(N)$.

3. Sparse extended form: the main idea. Our main observation is that Algorithm 3 can be rewritten as a multiplication of a sparse matrix by vector. The first step of Algorithm 3 can be rewritten as a matrix-vector product $Dx = \hat{x}_l$, where

$$D = \begin{bmatrix} D_1 & 0 & 0 & 0 \\ 0 & D_2 & 0 & 0 \\ 0 & 0 & \ddots & 0 \\ 0 & 0 & 0 & D_{N_{lr}} \end{bmatrix}, \tag{3.1}$$

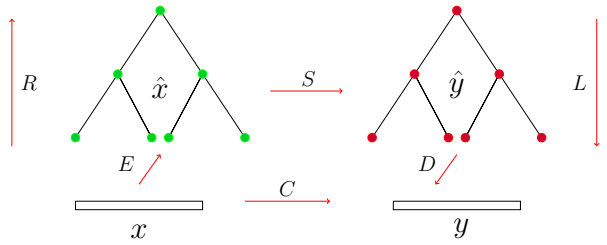

Fig. 2.2: Illustration of the Algorithm of matrix-vector product.

where N_{lr} is number of leaf nodes in tree $\mathcal{T}_r$, zeros represent zero matrices of appropriate sizes. We can rewrite the Step 2 as $R\hat{x} = \hat{x}_n$, where

$$R = \begin{bmatrix} R_1 & 0 & 0 & 0 \\ 0 & R_2 & 0 & 0 \\ 0 & 0 & \ddots & 0 \\ 0 & 0 & 0 & R_{N_{nr}} \end{bmatrix}, \tag{3.2}$$

where N_{nr} is the number of non-leaf nodes in the tree $\mathcal{T}_r$ zeros are zero matrices of appropriate sizes. The third step of Algorithm 3 corresponds to equation $L\hat{y}_n + S\hat{x} = \hat{y}$, where

$$L = \begin{bmatrix} L_1 & 0 & 0 & 0 \\ 0 & L_2 & 0 & 0 \\ 0 & 0 & \ddots & 0 \\ 0 & 0 & 0 & L_{N_{lc}} \end{bmatrix}, S = \begin{bmatrix} S_{11} & S_{12} & \cdots & S_{1N_c} \\ S_{21} & S_{22} & \cdots & S_{2N_c} \\ \vdots & \vdots & \ddots & \\ S_{N_c1} & S_{N_c2} & & S_{N_cN_c} \end{bmatrix}, \tag{3.3}$$

where N_c is the number of nodes in tree $\mathcal{T}_c$, N_{lc} is the number of leaf nodes in tree $\mathcal{T}_c$, $S_{ij} = 0$ if $i \in \mathcal{T}_c$ $j \in \mathcal{T}_r$ $(i,j) \in \mathcal{P}_{far}$. And the final step we rewrite as $y = Ey_l + Cx$, where

$$E = \begin{bmatrix} E_1 & 0 & 0 & 0 \\ 0 & E_2 & 0 & 0 \\ 0 & 0 & \ddots & 0 \\ 0 & 0 & 0 & E_{N_{lc}} \end{bmatrix}, \tag{3.4}$$

Putting this all together we get

$$\begin{cases} Dx = \hat{x}_l \\ R\hat{x} = \hat{x}_n \\ L\hat{y}_n + S\hat{x} = \hat{y} \\ E\hat{y}_l + Cx = y \end{cases}, \tag{3.5}$$

or in the block form:

$$\begin{bmatrix} C & 0 & 0 & E \\ 0 & S & L & 0 \\ 0 & R & 0 & 0 \\ D & 0 & 0 & 0 \end{bmatrix} \begin{bmatrix} x \\ \hat{x} \\ \hat{y}_n \\ \hat{y}_l \end{bmatrix} = \begin{bmatrix} y \\ \hat{y} \\ \hat{x}_n \\ \hat{x}_l \end{bmatrix}, \tag{3.6}$$

Denote the obtained matrix by H_0, and also introduce a block vector $\hat{x} = \begin{bmatrix} \hat{x}_n \\ \hat{x}_l \end{bmatrix}$ and $\hat{y} = \begin{bmatrix} \hat{y}_n \\ \hat{y}_l \end{bmatrix}$,. Finally, we get

$$H_0 \begin{bmatrix} x \\ \hat{x} \\ \hat{y} \end{bmatrix} = \begin{bmatrix} y \\ \hat{y} \\ \hat{x} \end{bmatrix}. \tag{3.7}$$

Recall that our goal is given y compute x, therefore

$$\left(H_0 + \begin{bmatrix} 0^{N \times N} & 0 & 0 \\ 0 & 0 & -I^{N_{\hat{y}} \times N_{\hat{y}}} \\ 0 & -I^{N_{\hat{x}} \times N_{\hat{x}}} & 0 \end{bmatrix} \right) \begin{bmatrix} x \\ \hat{x} \\ \hat{y} \end{bmatrix} = \begin{bmatrix} y \\ 0 \\ 0 \end{bmatrix},$$

where $N_{\hat{x}} = len(\hat{x})$, $N_{\hat{y}} = len(\hat{y})$ and

$$H = H_0 + \begin{bmatrix} 0 & 0 & 0 \\ 0 & 0 & -I \\ 0 & -I & 0 \end{bmatrix}.$$

The final system of equations has the form

$$H \begin{bmatrix} x \\ \hat{x} \\ \hat{y} \end{bmatrix} = \begin{bmatrix} y \\ 0 \\ 0 \end{bmatrix}. \tag{3.8}$$

Now the right hand side of system (3.8) contains only known values and we can solve it and find x, where x is the solution of $Ax = y$. We will call the matrix H "sparse extended form", or SE-form, of the $\mathcal{H}^2$-matrix A.

4. Properties of the SE-form. An important property of the SE-form is that if A non-singular, the SE-form of A is non-singular as well.

THEOREM 4.1. *If a matrix $A \in \mathbb{R}^{N \times N}$ is a nonsingular $\mathcal{H}^2$-matrix, then $H = \mathrm{SE}(A) \in \mathbb{R}^{N_H \times N_H}$ is nonsingular, and $N_H < (2k+1)N$, where k is the maximum of numbers of levels of the cluster trees $\mathcal{T}_c$ and $\mathcal{T}_r$.*

Proof. First let us prove that the matrix H is square. The matrix A is square therefore $\mathrm{len}(x) = \mathrm{len}(y) = N$. Let N_r be the number of rows in H and N_c be the number of columns in it. Then,

$$N_r = \mathrm{len}(x) + \mathrm{len}(\hat{x}) + \mathrm{len}(\hat{y}) = \mathrm{len}(y) + \mathrm{len}(\hat{y}) + \mathrm{len}(\hat{x}) = N_c$$

Thus H is a square matrix. Note $\mathrm{len}(\hat{x}) \leq k_1 N$, $\mathrm{len}(\hat{y}) \leq k_2 N$, where k_1 and k_2 are the numbers of levels of the trees $\mathcal{T}_c$ and $\mathcal{T}_r$. Therefore

$$N_H = \mathrm{len}(x) + \mathrm{len}(\hat{x}) + \mathrm{len}(\hat{y}) = N + k_1 N + k_2 N < (2k+1)N$$

Now we prove that H is nonsingular. Suppose that $Hz = 0$ and let us prove that it implies $z = 0$. Due to the construction of the SE-form, the first block component of the vector z satisfies $Ax = 0$, thus $x = 0$. According to Algorithm 3 and (3.5) $x = 0$ leads to $\hat{x} = 0$ and $\hat{y} = 0$ therefore $z = 0$ and the kernel of H is trivial. □ Note, that the condition number of the SE-form is typically much larger than of the original matrix, so special preconditioning is need

5. Solvers based on SE-form. How we can use the SE-form of the matrix for the solution of linear systems. We propose several methods, which are listed below.

Method 1: (Direct solver) Apply sparse direct solvers to $\mathrm{SE}(A)$ and given y compute x.

Method 2: (Preconditioning (3.8) with matrix $\mathrm{SE}(A)$) Construct preconditioner to $\mathrm{SE}(A)$ based on the block structure.

Method 3: (Iterative solvers for systems with A using $\mathrm{SE}(A)$ as preconditioner) SE-form can be used as preconditioner for the original system. To solve the correction equation we apply several steps of some solver for $\mathrm{SE}(A)$.

Now we describe them in more details.

5.1. Method 1. Applying any fast sparse direct solver to $\mathrm{SE}(A)$ is natural idea. However, it is ok for small N, but the amount of memory required for the such solver grows very quickly. The advantage of this method is that it is very simple to implement. Once a sparse representation of the $\mathrm{SE}(A)$ is given, we only need to call the procedure. Another approach is to compute some preconditioner for $\mathrm{SE}(A)$. In our numerical experiments we have tested ILUT preconditioner. We will call this approach SE-ILUT preconditioner.

5.2. Method 2. Now we consider system (3.8) with sparse extended matrix $\mathrm{SE}(A)$ and find solution of this system iteratively. In numerical experiments we found that $\mathrm{SE}(A)$ always have very large condition numbers. Thus, a preconditioner is needed. A natural way is to use the block structure of $\mathrm{SE}(A)$ to construct the preconditioner. We propose a **SE-Block preconditioner**. We compute an approximate inverse of the "far block" of $\mathrm{SE}(A)$ and all others are replaced by identity matrices:

$$B = \begin{bmatrix} I & 0 & 0 \\ 0 & P(S) & 0 \\ 0 & 0 & I \end{bmatrix}, \tag{5.1}$$

where $P(S)$ is some preconditioner for the block S. Note, that block S can be rectangular, in this case we construct preconditioner for the smallest square block that contains S. In experiments we have seen that the ill-conditioning of S is the reason why H is ill-conditioned.

5.3. Method 3. The matrix A can be considered as a Schur complement of $H = \mathrm{SE}(A)$ with components related to $\widehat{x}$ and $\widehat{y}$ eliminated. Typically, Schur complement is used as a preconditioner; here we use *reverse Schur component preconditioning* (similar ideas were used in [2]). In this method, Schur complement is used in the opposite fashion: we solve the system with a matrix A and to solve the correction equation

we go to the large system (extend the residual), apply several steps of some preconditioned iterative solver for SE(A) and extract the required vector as a corresponding component of the result. This approach appears to be the most effective one for the problems we have considered. Moreover, additional speedup can be obtained by using *recompression* of the $\mathcal{H}^2$-matrix, i.e.

$$A \approx B$$

where B has smaller ranks in the rank distribution. The description of the robust algorithm based on SVD can be found in [4], and it is implemented in the h2tools package as well. Then we use SE(B) instead of SE(A) in the algorithm, so this method is not a "reverse Schur complement", but "approximate reverse Schur complement" method. We will denote this approach *SVD-SE*. The final algorithm is summarized in Algorithm 4.

Algorithm 4: SVD-SE method

Data: Matrix A in the $\mathcal{H}^2$-format, right-hand side y, required tolerance ε, inner parameters: δ_{ILUT} and δ_{SVD}, number of reverse Schur iterations k_{schur}

Result: Approximate solution x such that $\|Ax - y\| \le \varepsilon$

begin

1. Compute B as a recompressed $\mathcal{H}^2$ representation of A with accuracy δ_{SVD}
2. Compute SE(B) and ILUT preconditioner P with drop tolerance δ_{ILUT}
3. $x = \text{GMRES}(A, y, \text{tol} = \varepsilon, \text{prec} = \text{RevSchur})$

Procedure `RevSchur(r)`

1. Extend right-hand side: $\widehat{r} = \begin{pmatrix} r^\top & 0 & 0 \end{pmatrix}^\top$
2. Do k_{schur} steps of ILUT-preconditioned GMRES for SE(B): $\widehat{z} = \text{GMRES}(\text{SE}(B), \widehat{r}, \text{maxit} = k_{\text{schur}}, \text{prec} = P)$
3. Extract z as the first N components of z
4. **return** z

6. Numerical experiments.

6.1. Electrostatic problem. As a model problem we consider boundary integral equation of the form

$$\int_\Omega \frac{q(y)}{\|x - y\|} dy = f(x), \quad x \in \Omega, \tag{6.1}$$

where Ω is $[0, 1]^2$. The function $f(x)$ is given and $q(y)$ is to be found. Equation (6.1) is discretized using collocation method with piece-wise constant basis functions on the triangular mesh Ω_N with N triangles (see Figure 6.2) The matrix elements can be evaluated analytically [14]. The matrix A is then approximated in the $\mathcal{H}^2$-format using the h2tools package [17]. The SE-form of A is presented on Figure 6.1 for $N = 1196$.

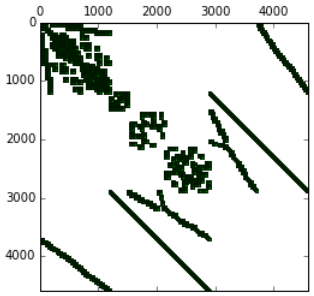

Fig. 6.1: SE-form of matrix A

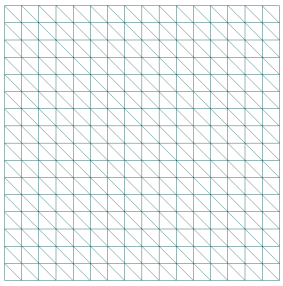

Fig. 6.2: Triangular mesh Ω_N

6.1.1. Method 1. In Table 6.3 the results of Method 1 (sparse direct solver applied to SE(A)) are presented. As it is readily seen, the memory quickly becomes a bottleneck.

N	time, (s)	Mem, (MB)
3928	2.585	376.85
13640	37.76	2527.7
28080	234	5012.1
59428	—	—
98936	—	—

Fig. 6.3: Timings and memory for Method 1

6.1.2. Method 2. The convergence of GMRES method with SE-ILUt and block preconditioners is presented on Figure 6.4.

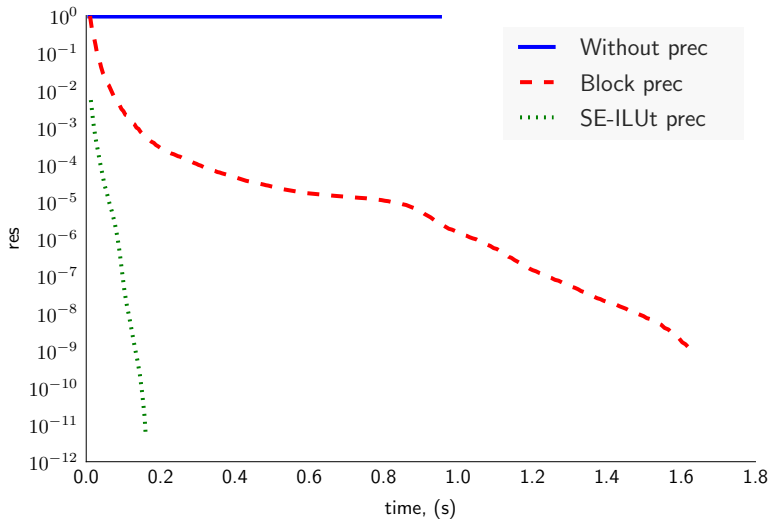

Fig. 6.4: Method 2: GMRES convergence with different preconditioners

The number of iterations with the SE-ILUT preconditioner is significantly smaller than the number of iterations with the block preconditioner, however, the computation of the block preconditioner is much less expensive. This is illustrated in Table 6.5.

N	Block prec + GMRES, (s)	SE-ILUt prec + GMRES, (s)
3928	0.085 + 1.28	1.75 + 0.17
13640	0.23 + 5.6	9.17 + 0.52
28080	0.53 + 11.8	27.17 + 0.91
59428	1.34 + 34.8	75.02+ 3.13
98936	3.28 + 59.13	134.11+ 10.2

Fig. 6.5: Timing for building the preconditioner and solving the system using the GMRES method

6.1.3. Method 3. The convergence of the GMRES method with reverse-Schur preconditioning is presented on Figure 6.6 for $N = 28080$.

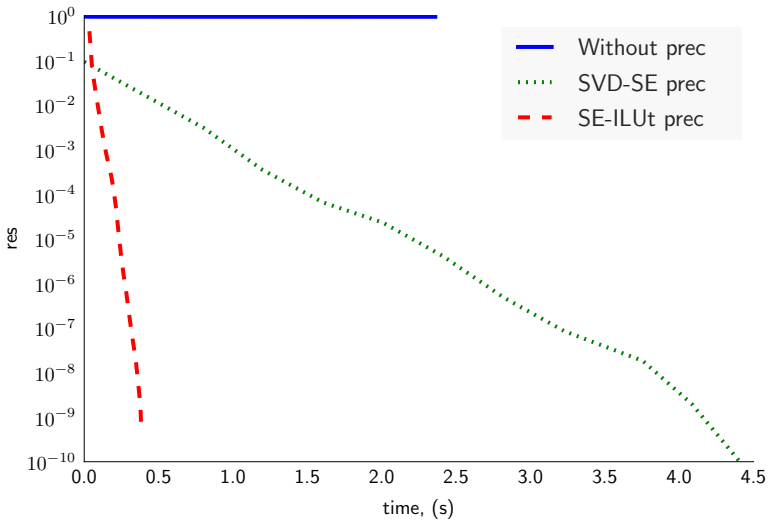

Fig. 6.6: Method 3: Convergence of GMRES for $N = 28080$

The total computational cost for different reverse-Schur preconditioner is given in Table 6.1. Note that the SE-ILUt preconditioner in the second column is the same as in the previous section, however it is more effective to use it as a reverse-Schur preconditioner for the original system, than for the full SE-form. The compression of the $\mathcal{H}^2$-form of the matrix yields the best preconditioner by a factor of 4.

N	SE-ILUt prec + GMRES	SVD-SE prec + GMRES, (s)
3928	2.13 + 0.03	0.21 + 0.11
13640	8.84 + 0.11	1.34 + 1.32
28080	24.8 + 0.35	8.35 + 2.94
59428	69.1 + 1.33	19.7 + 6.13
98936	150.2 + 4.38	40.01 + 15.03

Table 6.1: Timings for the Method 3 with different reverse-Schur preconditioners

6.1.4. Final comparison. In Table 6.2 we present final comparison for all methods for different N. For Method 2 and Method 3 the solution time is the total computational time (building the preconditioner and solving the system using GMRES)

	Method 1	Method 2		Method 3	
N	Dir. sol, (s)	Block, (s)	SE-ILUt, (s)	SE-ILUt, (s)	SVD-SE, (s)
3928	2.585	3.215	2.11	2.16	0.87
13640	37.76	10.83	9.69	8.95	4.65
28080	234	22.33	28.08	25.47	16.92
59428	—	53.14	78.15	70.37	42.01
98936	—	122,41	144.31	127.95	89.94

Table 6.2: Timings for all methods for different N.

The memory is an important constraint, and in Table 6.7 we present the memory required for each of the methods for different N.

	Method 1	Method 2		Method 3	
N	Dir. sol, (MB)	Block, (MB)	SE-ILUt, (MB)	SE-ILUt, (MB)	SVD-SE, (MB)
3928	424	24.2	31.4	15.7	< 10
13640	2606	251.2	279.8	31.4	< 10
28080	14702	675.3	2464.9	219.8	47.1
59428	—	1301	9720.6	810	128.3
98936	—	6340.1	—	2028.8	221.5

Fig. 6.7: Memory requirements for all the methods for different N, missing entries mean "out of memory"

Method 3 with SVD recomporession is the fastest method and requires much less memory.

6.2. Hypersingular integral equation. The second problem is the hypersingular integral equation

$$\int_{\Omega} \frac{q(y)}{\|x-y\|^3}\, dy = f(x), \tag{6.2}$$

where Ω is shown on Figure 6.9. The equation is discretized using the collocation method with piecewise-constant basis functions (also known as discrete vortices method), and the approximation in the $\mathcal{H}^2$-format is done using the h2tools package as well. The SE-form of A is given on Figure 6.8. The mesh has $N = 28340$ triangles.

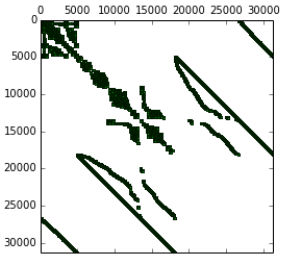

Fig. 6.8: SE-form of matrix A

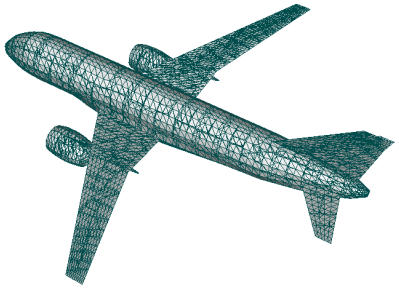

Fig. 6.9: Surface Ω_N

The comparison of different methods is shown in Table 6.10.

	Method 1	Method 2		Method 3	
	Dir. sol	Block	SE-ILUt	SE-ILUt	SVD-SE
memory, (M)	-	28.3	31.7	31.7	17.3
time, (s)	-	-	-	8.28	13.4

Fig. 6.10: Timing and memory for all methods for problem (6.2)

7. Conclusions. The new SE-form of the $\mathcal{H}^2$-matrix allows for different ways for the constuction of effective linear systems solvers with such matrices. Numerical experiments show that the most effective way is to use the reverse-Schur preconditioner with SVD-recompression of the $\mathcal{H}^2$-form of A. The implementation of the methods is available as a part of the open-source package h2tools [17], and the numerical experiments are available as IPython notebooks.

REFERENCES

[1] S. Ambikasaran and E. Darve, *The inverse fast multipole method*, arXiv preprint, 2014.
[2] J. Bardhan, M. Altman, B. Tidor, and J. White, *"Reverse-Schur" approach to optimization with linear PDE constraints: Application to biomolecule analysis and design*, J. Chem. Theory Comput., 5 (2009), pp. 3260–3278.
[3] M. Bebendorf, *Hierarchical LU decomposition-based preconditioners for BEM*, Computing, 74 (2005), pp. 225–247.
[4] S. Börm, *Efficient numerical methods for non-local operators: H2-matrix compression, algorithms and analysis*, vol. 14, European Mathematical Society, 2010.
[5] S. Borm, L. Grasedyck, and W. Hackbusch, *Introduction to hierarchical matrices with applications*, Eng. Anal. Bound Elem., 27 (2003), pp. 405–422.
[6] S. Chandrasekaran, P. Dewilde, M. Gu, W. Lyons, and T. Pals, *A fast solver for HSS representations via sparse matrices*, SIAM J. Matrix Anal. Appl., 29 (2006), pp. 67–81.

[7] S. Chandrasekaran, M. Gu, X. Li, and J. Xia, *Some fast algorithms for hierarchically semiseparable matrices*, Private Communication, (2007).

[8] S. Chandrasekaran, M. Gu, and W. Lyons, *A fast adaptive solver for hierarchically semiseparable representations*, Calcolo, 42 (2005), pp. 171–185.

[9] S. Chandrasekaran, M. Gu, and T. Pals, *A fast ulv decomposition solver for hierarchically semiseparable representations*, SIAM J. Matrix Anal. A., 28 (2006), pp. 603–622.

[10] E. Corona, P.-G. Martinsson, and D. Zorin, *An $O(N)$ direct solver for integral equations on the plane*, Appl. Comput. Harmon. A., (2014).

[11] L. Greengard and V. Rokhlin, *A fast algorithm for particle simulations*, J. Comput. Phys., 73 (1987), pp. 325–348.

[12] L. Grengard and V. Rokhlin, *The rapid evaluation of potential fields in three dimensions*, Springer, 1988.

[13] W. Hackbusch and S. Borm, *$\mathcal{H}^2$-matrix approximation of integral operators by interpolation*, Appl. Numer. Math., 43 (2002), pp. 129–143.

[14] J. L. Hess and A. Smith, *Calculation of non-lifting potential flow about arbitrary three-dimensional bodies*, tech. rep., DTIC Document, 1962.

[15] K. Ho and L. Greengard, *A fast direct solver for structured linear systems by recursive skeletonization*, SIAM J. Sci. Comput., 34 (2012), pp. A2507–A2532.

[16] P. G. Martinsson, *A fast randomized algorithm for computing a hierarchically semiseparable representation of a matrix*, SIAM J. Matrix Anal. A., 32 (2011), pp. 1251–1274.

[17] A. Mikhalev. `https://bitbucket.org/muxas/h2tools`, 2013.

[18] A. Mikhalev and I. V. Oseledets, *Adaptive nested cross approximation of non-local operators*, arXiv preprint 1407.1572, 2013.

[19] V. Rokhlin, *Rapid solution of integral equations of classical potential theory*, J. Comput. Phys., 60 (1985), pp. 187–207.

[20] H. Samet, *The quadtree and related hierarchical data structures*, ACM Comput. Surv., 16 (1984), pp. 187–260.

[21] S. D. Senturia, R. M. Harris, B. P. Johnson, S. Kim, K. Nabors, M. A. Shulman, and J. K. White, *A computer-aided design system for microelectromechanical systems (memcad)*, Mater. Res. Soc. Symp. P., 1 (1992), pp. 3–13.

[22] Z. Sheng, P. Dewilde, and S. Chandrasekaran, *Algorithms to solve hierarchically semiseparable systems*, in System theory, the Schur algorithm and multidimensional analysis, Springer, 2007, pp. 255–294.

[23] E. E. Tyrtyshnikov, *Mosaic-skeleton approximations*, Calcolo, 33 (1996), pp. 47–57.

[24] J. Xia, S. Chandrasekaran, M. Gu, and X. S. Li, *Superfast multifrontal method for large structured linear systems of equations*, SIAM J. Matrix Anal. A., 31 (2009), pp. 1382–1411.